

CLAIMS

1. An integrated circuit device, comprising:
 - a processing component;
 - a cache, which is arranged to store data for use by
 - 5 the processing component responsively to an addressing scheme based on memory addresses having an address length of m_1 bits; and
 - first and second buses coupled between the processing component and the cache, the buses having bus
 - 10 widths of n_1 and n_2 bits, respectively, such that $n_1 < m_1$,
 - the processing component and the cache each comprising a respective address bus expander coupled to the first bus in order to compact at least some of the memory addresses for transmission over the first bus so
 - 15 that each of the at least some memory addresses is transmitted over the first bus in one cycle of the first bus.
2. The device according to claim 1, wherein the data
- 20 comprise data words having a word length of m_2 bits stored at each address, such that $n_2 < m_2$, and
- wherein each of the processing component and the cache further comprises a respective second bus expander coupled to the second bus in order to compact at least
- 25 some of the data words for transmission over the second bus so that each of the at least some of the data words is transmitted over the second bus in one cycle of the second bus.
3. The device according to claim 2, wherein the data
- 30 words comprise data values for processing by the device, and wherein the processing component is arranged to load

the compacted data words via the second bus from the cache for processing and to store the compacted data words via the second bus to the cache.

5 4. The device according to claim 2, wherein the data words comprise instructions for execution by the device, wherein the compacted words comprise compacted instructions, and wherein the processing component is arranged to fetch the compacted instructions via the
10 second bus.

5. The device according to claim 2, wherein the address bus expander and the second bus expander are arranged to compact the memory addresses and the data words
15 simultaneously, so as to transmit a compacted memory address and a compacted data word for storage at the memory address together in one cycle of the first and second buses.

20 6. The device according to claim 2, wherein the address bus expander and the second bus expander are arranged to compact the memory addresses and the data words by transmitting indices to values in respective tables held by the bus expanders, and wherein the cache is arranged
25 to store at least some of the indices together with the data.

7. The device according to claim 1, wherein the address bus expander is arranged to compact each of the at least
30 some of the memory addresses by dividing each of the

memory addresses into at least first and second fields,
storing values of the second field in a respective table
such that the values in respective tables held by the
address bus expander in the processing component and the
5 address bus expander in the cache are identical, and if
the second field of a memory address matches a value in
the table, transmitting an index corresponding to the
value over the first bus along with the first field in
the one cycle of the bus.

10

8. The device according to claim 7, wherein the first
field comprises a set of least significant bits (LSB) of
the memory address, while the second field comprises a
set of most significant bits (MSB) of the memory address.

15

9. The device according to claim 7, wherein the at
least first and second fields comprise a third field, and
wherein the address bus expander is arranged to compact
each of the at least some of the memory addresses by
20 transmitting first and second indices corresponding to
the values of the first and third fields, respectively,
over the first bus along with the first field.

10. The device according to claim 7, wherein the address
25 bus expander in the processing component is arranged,
when the second field of the memory address does not
match any of the values in the table, to transmit both of
the first and second fields over multiple cycles of the
bus, and to cause the respective table of the bus
30 expander to be updated in both the processing component
and the cache.

11. The device according to claim 7, wherein the cache comprises lines of the data that are indexed according to the first field, each line containing a corresponding
5 value of the second field, and

wherein the address bus expander in the cache is arranged, upon receiving the index over the first bus, to retrieve the value of the second field from the table responsively to the index, and

10 wherein the cache is arranged to determine whether a cache hit has occurred by checking the retrieved value against the corresponding value of the second field in the line that is indexed by the first field.

15 12. The device according to claim 11, wherein the address bus expander is arranged to retrieve the value of the second field from the table simultaneously with retrieval of the data from the line in the cache that is indexed by the first field for transmission of the data
20 over the second bus to the processing component.

13. The device according to claim 1, wherein the cache that is coupled to the processing component by the first and second buses is a Level 1 (L1) cache, and further
25 comprising a Level 2 (L2) cache, and third and fourth buses coupling the L2 cache to the L1 cache, the L1 cache and the L2 cache comprising further bus expanders coupled to at least one of the third and fourth buses.

14. A method for coupling a processing component to a cache in an integrated circuit device, the method comprising:

5 configuring the cache to store data for use by the processing component responsively to an addressing scheme based on memory addresses having an address length of m_1 bits;

10 coupling first and second buses coupled between the processing component and the cache, the buses having bus widths of n_1 and n_2 bits, respectively, such that $n_1 < m_1$;

compacting at least some of the memory addresses for transmission over the first bus so that each of the at least some memory addresses can be transmitted over the first bus in one cycle of the first bus;

15 transmitting at least the compacted memory addresses over the first bus; and

conveying the data over the second bus responsively to the compacted memory addresses.

20 15. The method according to claim 14, wherein the data comprise data words having a word length of m_2 bits stored at each address, such that $n_2 < m_2$, and

25 wherein conveying the data comprises compacting at least some of the data words for transmission over the second bus so that each of the at least some of the data words is transmitted over the second bus in one cycle of the second bus.

30 16. The method according to claim 15, wherein the data words comprise data values for processing by the integrated circuit device, and wherein conveying the data

comprises loading the compacted data words via the second bus from the cache for processing by the processing component, and storing the compacted data words from the processing component via the second bus to the cache.

5

17. The method according to claim 15, wherein the data words comprise instructions for execution by the device, wherein the compacted words comprise compacted instructions, and wherein conveying the data comprises
10 fetching the compacted instructions from the cache via the second bus to the processing component.

18. The method according to claim 15, wherein compacting the at least some of the memory addresses and compacting
15 at least some of the data words comprise compacting at least a portion of the memory addresses and the data words simultaneously, so as to transmit a compacted memory address and a compacted data word for storage at the memory address together in one cycle of the first and
20 second buses.

19. The method according to claim 15, wherein compacting the at least some of the memory addresses and compacting at least some of the data words comprise compacting the
25 memory addresses and the data words by transmitting indices to values in respective tables held in the cache and the processing component, and storing at least some of the indices in the cache together with the data.

20. The method according to claim 14, wherein compacting the at least some of the memory addresses comprises:

dividing each of the at least some of the memory addresses into at least first and second fields; and

5 storing values of the second field in a respective table such that the values in respective tables held in the processing component and in the cache are identical, and

wherein transmitting at least the compacted memory
10 addresses comprises, if the second field of a memory address matches a value in the table, transmitting an index corresponding to the value over the first bus along with the first field in the one cycle of the bus.

15 21. The method according to claim 20, wherein the first field comprises a set of least significant bits (LSB) of the memory address, while the second field comprises a set of most significant bits (MSB) of the memory address.

20 22. The method according to claim 20, wherein the at least first and second fields comprise a third field, and wherein transmitting the index comprises transmitting first and second indices corresponding to the values of the first and third fields, respectively, over the first
25 bus along with the first field.

23. The method according to claim 20, wherein transmitting at least the compacted memory addresses comprises, when the second field of the memory address
30 does not match any of the values in the table, transmitting both of the first and second fields over

multiple cycles of the bus, and updating the respective table in both the processing component and the cache.

24. The method according to claim 20, wherein
5 configuring the cache comprises indexing lines of the data that according to the first field, each line containing a corresponding value of the second field, and wherein conveying the data comprises:

receiving the index in one of the compacted memory
10 addresses over the first bus; and

determining whether a cache hit has occurred by
retrieving the value of the second field from the table
responsively to the index, and checking the retrieved
value against the corresponding value of the second field
15 in the line that is indexed by the first field.

25. The method according to claim 24, wherein
determining whether the cache hit has occurred comprises
simultaneously retrieving the value of the second field
20 from the table and retrieving the data from the line that
is indexed by the first field for transmission over the
second bus to the processing component.

26. The method according to claim 14, wherein the cache
25 that is coupled to the processing component by the first
and second buses is a Level 1 (L1) cache, and further
comprising coupling a Level 2 (L2) cache via third and
fourth buses to the L1 cache, and compacting at least
some of the memory addresses or the data for transmission
30 over at least one of the third and fourth buses.